

**APPLICATION FOR  
UNITED STATES PATENT**

**in the names of**

**Keren Perlmutter, Sharon Perlmutter,  
Eric Wang, and Paul Klamer**

**for**

**REDUCTION OF DIFFERENTIAL  
RESOLUTION OF SEPARATIONS**

Fish & Richardson P.C.  
1425 K Street, NW  
11<sup>th</sup> Floor  
Washington, DC 20005-3500  
Tel.: (202) 783-5070  
Fax: (202) 783-2331

**ATTORNEY DOCKET:  
06975-366001**

# **REDUCTION OF DIFFERENTIAL RESOLUTION OF SEPARATIONS**

## **CROSS-REFERENCE TO RELATED APPLICATIONS**

This application claims priority from and is a continuation-in-part of U.S. Application No. 10/035,337, filed January 4, 2002, and titled “Registration of Separations,” which is incorporated by reference. This application also claims priority from U.S. Provisional Application No. 60/434,650, filed December 20, 2002, and titled 5 “Reduction of Differential Resolution of Separations,” which is incorporated by reference.

## **TECHNICAL FIELD**

Certain disclosed systems relate generally to image processing, and more particularly to the reduction of distortion that is manifested from separations with 10 different resolutions.

## **BACKGROUND**

Color motion picture film is a relatively recent development. Before the advent of color film stock in the 1950s, a process for making color motion pictures included capturing color information on two or more reels of black and white film. In the original 15 Technicolor three-color film separation process, three reels of black and white film were loaded into a specially-designed movie camera. The light coming through the lens was split into the three primary colors of light and each was recorded on a separate reel of black and white film. After developing the three reels, three photographic negatives were created representing the yellow (inverted blue), the cyan (inverted red), and the magenta 20 (inverted green) portions of the original scenes.

In addition to the creation of color separations through the original Technicolor process, color separations also have been produced and used for the archival of color film because black and white film stock generally has a much greater shelf-life than color film. In this process, the color film stock is used to expose one reel of black and white 5 film with sequential records of red, green, and blue so that each frame is printed three times on the resultant reel to form a sequential separation.

Film studios may recombine the three color separations onto a single reel of color film using a photographic process that is performed in a film laboratory. In the case of three color separations that are each located on a separate reel, an optical film printer is 10 employed to resize and reposition each source reel, one at a time. In particular, three passes are made. First, the magenta source reel is projected through an appropriate color filter onto the destination reel. Thereafter, the destination reel is rewound, the next source reel is loaded and resized, and the color filter is changed. The process is repeated until all three color separations have been printed on the single destination reel using the 15 optical film printer. The resulting destination reel is called an interpositive (“IP”), and the colors are now represented as red, green, and blue (as opposed to cyan, magenta, and yellow).

The Technicolor three-color film separation process, as well as other processes, is subject to a variety of film distortions, including, for example, differential resolution. 20 Differential resolution may arise because, for example, the nature of the light path and lens coatings in the Technicolor cameras typically cause the three film separations to have drastically different resolution or sharpness. The cyan filter typically is located behind the yellow filter in what is known as a bipack arrangement. Light that passes

through the yellow filter is filtered and, unfortunately, diffused before striking the cyan filter. As a result, the yellow (inverted blue) separation typically has a greater resolution compared to the cyan (inverted red) separation. The magenta (inverted green) separation is not created with a bipack arrangement and typically has a resolution that is similar to 5 that of the yellow (inverted blue) separation. This difference in resolution may result in red fringing or blurring around edges in the picture or image.

## SUMMARY

An implementation described below addresses the lower resolution of the red separation by increasing the resolution of selected portions of the red separation. The 10 portions are selected using information from a higher resolution separation (blue or green). For example, such information may include determining that a particular edge in the higher resolution separation corresponds to an edge in the red separation, in which case that particular edge in the red separation may become a selected portion. After selecting the portions, the resolution of the selected portions is increased by using 15 information produced by the application of a wavelet transformation to the higher resolution separation to modify corresponding information produced by the application of a wavelet transformation to the red separation. For example, various coefficients produced by the application of the wavelet transformation to the higher resolution separation (or a function of these coefficients) may be copied to a set of coefficients 20 produced by the application of the wavelet transformation to the red separation, where such coefficients impact the selected portions.

This implementation provides an automatic and efficient differential resolution reduction process for color film separations. Further, the process requires minimal human intervention and determines where differential resolution should be reduced or corrected within an image.

5        Many implementations may be characterized as including a “where” operation and a “how” operation. The “where” operation determines where to modify an image, and may do so, for example, by determining the portion(s) at which one or more properties of an image are to be modified. The “how” operation determines how to modify the portion(s) of the image identified in the “where” operation. Either or both of 10 the “where” and “how” operations may use, for example, frequency-based information, time-based information, or both, and the information may be, for example, intra-frame or inter-frame.

According to one aspect, reducing differential resolution includes selecting a first image containing first information about a scene and selecting a second image containing 15 second information about the scene. A portion of the first image and a portion of the second image have differential resolution. A location at which to modify a property of the first image to reduce the differential resolution is determined, with the location being in the portion of the first image and the determination being based on information obtained at least in part from the portion of the second image. The differential resolution 20 is reduced by modifying the property at the determined location in the portion of the first image.

The first image and the second image may be digital images. In addition, the location may include a pixel, and the property may include an intensity value of the pixel or a function of the intensity value.

The first image and the second image may include color separations of a film  
5 frame or may be extracted from a composite color image. The color composite image may be generated from color separations of a film frame. The first image may include a red separation and the differential resolution may result in red fringing. A non-modify location at which the property is not to be modified may be determined in the first image.

Determining the location may include comparing one or more features of an edge  
10 in the first image with one or more features of an edge in the second image. The edge may be selected as a location to modify based on a result of the comparison.

Determining the location also may include selecting one or more edges to modify. Selecting one or more edges to modify may include selecting, for one of the one or more edges, a single edge pixel of an edge that includes multiple edge pixels. Selecting one or  
15 more edges to modify may include (i) comparing one or more features of an edge in the first image with one or more features of an edge in the second image, and (ii) selecting the edge as an edge to modify based on a result of the comparison. The one or more features may include a feature selected from the group consisting of a location of the edge, a direction of the edge, an extent of the edge, an intensity-change direction, and an  
20 intensity range traversed. An edge extent to modify may be determined for each selected edge.

Multiple edges may be selected for modification. Two of the selected edges may be connected based on properties of the two selected edges, and an edge extent may be

determined for the connected selected edges. Connecting two selected edges may be based on spatial proximity between the two selected edges or one or more of intensity differences between particular pixels in each of the two selected edges and intensity differences between particular pixels spatially located between the two selected edges.

- 5 Determining an edge extent for the connected selected edges may be based on edge extents that would have been determined for each of the selected edges before the edges were connected.

A selected edge may be unselected based on the size of the selected edge.

- 10 Determining the location in the portion of the first image may be based on information obtained at least in part from the portion of the second image. The information may be for a first direction only. Modifying the property at the location may include producing a modified first image. A location may be determined at which to modify the property in the modified first image. The determination may be based on information obtained at least in part from the second image. The information obtained from the second image
- 15 may be for a second direction that is orthogonal to the first direction. The property may be modified at the location in the modified first image.

- 20 The first image may include an image that has been modified with information obtained from the second image. Selecting a second image may include selecting a second image from a set of images based on one or more criteria. The one or more criteria may include intensity information and resolution information.

Determining the location may be performed automatically or interactively. A feathering technique may be applied to a region of the first image that includes the

location to modify. The feathering technique may be applied after the location is modified.

Modifying the property at the location in the first image may include applying a first wavelet transformation to the portion of the first image to produce a result, and 5 applying a second wavelet transformation to the portion of the second image. One or more coefficients produced by the application of the first wavelet transformation may be modified based on one or more coefficients produced by the application of the second wavelet transformation to produce a modified result.

A non-modify location may be determined in the first image at which the property 10 is not to be modified. An inverse wavelet transformation may be applied to the modified result of the first wavelet transformation to produce a digital image. When the property has been modified at the non-modify location in the digital image, the property at the non-modify location may be restored to its original value.

An inverse wavelet transformation may be applied to the modified result, and a 15 determination may be made as to whether the differential resolution is reduced between the portion of the first image and the portion of the second image. An inverse wavelet transformation may be applied to the modified result to produce another result, and a feathering technique may be applied to a portion of the other result that includes the determined location to modify. Applying the feathering technique may include linearly 20 interpolating between intensity values within the portion of the other result.

According to another aspect, selecting an edge includes accessing a first image and accessing a second image. The second image includes information that is complementary to information in the first image. A feature of an edge in the first image

is compared with a feature of a corresponding edge in the second image. The edge in the first image is selected, based on a result of the comparison, as an edge to modify.

The selected edge may be modified using resolution information about the corresponding edge in the second image. Modifying the selected edge may include 5 modifying information produced by application of a first wavelet transformation to a portion of the selected edge. The modification may be based on information produced by application of a second wavelet transformation to a portion of the corresponding edge.

The first wavelet transformation may be the same as the second wavelet transformation.

Modifying information produced by application of the first wavelet 10 transformation may include using one or more coefficients produced by application of the second wavelet transformation to replace one or more coefficients produced by application of the first wavelet transformation.

Using one or more coefficients produced by application of the second wavelet transformation may include scaling a coefficient produced by application of the second 15 wavelet transformation to produce a scaled coefficient. The scaled coefficient may be used to replace a coefficient produced by application of the first wavelet transformation.

According to another aspect, modifying a property of an image includes accessing a first image containing first information about a scene, and accessing a second image containing second information about the scene. A portion of the first image and a portion 20 of the second image have differential resolution. A location is determined at which to modify a property of the first image to reduce the differential resolution. The determining is based on a time-domain comparison of the portion of the first image and the portion of the second image. The property is modified at the location by modifying

information produced by application of a wavelet transformation to the portion of the first image. The information is modified based on information produced by application of a wavelet transformation to the portion of the second image.

An apparatus may include a computer readable medium on which are stored 5 instructions that, when executed by a machine, result in various of the above operations being performed. The apparatus may include a processing device coupled to the computer readable medium for executing the stored instructions.

One or more implementations are set forth in the accompanying drawings and the description below. Other implementations will be apparent from the description, the 10 drawings, and the claims.

## **DESCRIPTION OF DRAWINGS**

The patent or application file contains at least one drawing executed in color. Copies of this patent or patent application publication with color drawing(s) will be provided by the Office upon request and payment of the necessary fee.

15 FIG. 1A is an image in which red fringing is manifested from differential resolution of separations.

FIG. 1B is a color image of a portion of FIG. 1A, illustrating the red fringing.

FIG. 2 is the color image of FIG. 1B after being processed to reduce the differential resolution.

20 FIG. 3 is a block diagram of a system for reducing differential resolution.

FIG. 4 is a flow chart of a process to determine whether a pixel is to be modified.

FIG. 5 is an edge map for a simple low resolution digital image.

FIG. 6 is an edge map for a simple reference digital image.

FIG. 7 is an edge extent for the edges in FIG. 5.

FIG. 8 is an edge extent for the edges in FIG. 6.

FIG. 9 shows the edges from the edge map in FIG. 5 that have a matching edge in  
5 the edge map in FIG. 6.

FIG. 10 shows the edge map in FIG. 9 after applying a continuity operation.

FIG. 11 is a flow chart of a process for pruning “modify” pixels.

FIG. 12 is an edge extent for the edge map in FIG. 10.

FIG. 13 is a classification map showing “modify” pixels.

10 FIG. 14 is a flow chart of a process for modifying selected portions of a low  
resolution digital image.

FIG. 15 is a table showing the correlation between wavelet coefficients and pixels  
in a row of a digital image.

FIG. 16 is a map of a portion of a column of an image, providing certain  
15 information used in a feathering operation.

FIG. 17 is a classification map showing pixels that may be smoothed.

FIG. 18 is a block diagram of a system for implementing disclosed features.

**DETAILED DESCRIPTON**

FIG. 1A shows a portion of a scene in which red fringing occurs around various  
20 objects. FIG. 1B is a color image of a portion of FIG. 1A illustrating red fringing around  
edges along a pair of white pants 110 and along a boot 120. For example, a region 130  
shows red fringing to the right of boot 120. Also, a region 140 shows red fringing above

white pants 110. The red fringing is a result of differential resolution between the separations of the color image. As described above, the red fringing may have been introduced by the filming process.

FIG. 2 shows the same color image from FIG. 1B after the color separations are 5 processed using a technique described below to reduce the red fringing. As FIG. 2 illustrates, the red fringing around white pants 110, particularly in region 140 above the contoured edge of white pants 110, has been significantly reduced or eliminated. The red fringing in region 130 along boot 120 also has been significantly reduced or eliminated. Further, processing the image of FIG. 1B to reduce the red fringing has not visibly 10 degraded the resolution or quality of the image in other aspects.

FIG. 3 is a block diagram of a system 300 for reducing differential resolution. As an overview, system 300 performs the two general operations: (1) identifying locations for which the resolution should be modified and (2) modifying the resolution for those locations. The description below generally refers to individual images (frames), with the 15 understanding that the process is repeated for each image in a reel.

### Digitization Unit

System 300 includes a digitization unit 310 that receives the three separation images C (cyan), M (magenta), and Y (yellow), and provides three digital color 20 component images (“digital images”)  $R_D$  (red),  $G_D$  (green), and  $B_D$  (blue). The subscript D indicates that the image is digitized. Digitization unit 310 thus performs the steps of inverting the photographic negatives (C, M, Y) into photographic positives and digitizing the resulting positives into three digital images ( $R_D$ ,  $G_D$ ,  $B_D$ ).

Each digital image includes an array of pixels having a width and a height.

Within the array, each pixel location can be characterized by:

(x, y), where  $0 \leq x \leq \text{width}$ , and  $0 \leq y \leq \text{height}$ .

Each pixel location has an associated gray-level value that can be characterized by:

5  $I(x, y)$ , where  $0 \leq I(x, y) \leq I_{\text{max}}$ .

$I(x, y)$  represents the intensity of a particular color (for example, red, green, or blue) at pixel location (x, y).  $I_{\text{max}}$  represents the maximum possible intensity value for the pixel. For example, in the case of 16-bit data,  $I_{\text{max}} = 65,535$ .

10 Pre-Processing Unit

System 300 also includes a pre-processing unit 320 that receives digital images  $R_D$ ,  $G_D$ , and  $B_D$  from digitization unit 310 and produces modified digital images  $R$ ,  $G$ , and  $B$ . Pre-processing unit 320 is an optional unit in system 300 and performs pre-processing operations on one or more of the three digital images  $R_D$ ,  $G_D$ , and  $B_D$ . For example, one or more of the images  $R_D$ ,  $G_D$ , and  $B_D$  may be smoothed, may be altered with a registration algorithm so as to be better aligned with the other digital images that compose the frame, or may be processed with a grain reduction algorithm. In one implementation, pre-processing unit 320 operates on all three digital images  $R_D$ ,  $G_D$ , and  $B_D$ , and performs the following operations: (i) grain reduction using smoothing, (ii) registration, and (iii) an additional smoothing operation. If pre-processing unit 320 is not included in system 300, then images  $R$ ,  $G$ , and  $B$  will be the same as images  $R_D$ ,  $G_D$ , and  $B_D$ , respectively.

### Classification Unit

System 300 also includes a classification unit 330 that receives digital images R, G, and B from pre-processing unit 320 and provides a classification map for one of the images. For example, a classification map  $C_R$  identifies one or more locations (pixels) 5 within the red digital image for which the resolution is to be increased.  $C_R$  can be represented by the same array as the red digital image, except that each pixel contains either “modify” (M), “potentially modify” (PM), or “non-modify” (NM) labels rather than an intensity value. “Modify” indicates that the pixel intensity value of the corresponding pixel location in the red digital image is to be modified; “potentially 10 modify” indicates that the pixel intensity value might be modified, depending on, for example, further testing; and “non-modify” indicates that the pixel intensity value is not to be modified.

In one implementation, described in detail below, a label of PM is only an interim 15 label and the output of classification unit 330 is a classification map  $C_R$  in which the pixels are labeled either M or NM, but not PM. The use of the interim PM label allows the implementation to capture and to use the fact that a particular pixel may have passed at least some of the tests that indicate that a pixel should be modified.

Determining whether to modify a given pixel location can be based on a variety of factors. Because red fringing is a phenomenon associated with edges, classification unit 20 330 uses edge information to determine which pixels to modify. Edge information may be obtained for edges in one dimension or multiple dimensions. For example, classification unit 330 may use edge information pertaining to two-dimensional edges. Conversely, classification unit 330 may perform two iterations, with the first iteration

using edge information pertaining only to, for example, horizontal edges and the second iteration using edge information pertaining only to, for example, vertical edges.

Determining whether to modify a given pixel in the red digital image is based on information in one or more of the other digital images G and B, as well as on information 5 in the red digital image R. Classification unit 330 determines which of the other digital images to use to provide information for use in modifying the red digital image, and the digital image used to provide information is referred to as the reference digital image.

Various criteria may be used to select the reference digital image, including, for example, intensity value information and resolution information, as the following examples 10 indicate.

In one implementation, if the green and blue digital images both have an edge that matches an edge in the red digital image, the digital image (green or blue) that has the minimum average intensity at the ending (across one or more pixels) of the edge transition is selected as the reference digital image for that edge. In another

15 implementation, if the green and blue digital images both have an edge that matches an edge in the red digital image, the digital image (green or blue) that has the highest resolution is selected as the reference digital image. Resolution, or sharpness, can be defined in a variety of ways. For example, resolution for a set of pixels may be defined as a ratio of the range of intensity values present in the pixels (that is, the maximum 20 intensity value minus the minimum intensity value) over the number of pixels in the set.

In a further implementation, the selection criteria for selecting the reference digital image to be used for one particular edge may be different from the selection criteria used for another edge. When more than one reference digital image can be used, the classification

map  $C_R$  produced by classification unit 330 may provide not only M or NM information for each edge pixel, but also may provide information about which reference digital image is used for each edge pixel.

FIG. 4 is a flow chart of a process 400 used by one implementation of  
5 classification unit 330 to determine whether a pixel in the red digital image is to be modified. Process 400 includes obtaining a set of edges (an edge map) for R and one or more of G and B (410). The edge maps may be obtained by, for example, using an edge detection filter, such as, for example, a Canny filter. The pixels making up the edge are referred to as edge pixels or edge transition pixels. The edge information may be  
10 combined or may be separated. For example, the edge information may be separated into sets of information that capture edges in orthogonal directions. One implementation obtains separate horizontal and vertical edge information. FIGS. 5 and 6 illustrate edge maps for a simple red digital image and a simple reference digital image, respectively.

FIG. 5 shows a sample edge map 500 for a red digital image that has, for  
15 simplicity, only one-dimensional horizontal edges. Edge map 500 shows four edges. The four edges are: (i) a one-dimensional horizontal edge covering pixels from (1,3) to (3, 3), (ii) a one-dimensional horizontal edge covering pixels from (5,1) to (9,1), (iii) a one-dimensional horizontal edge covering pixels from (3,8) to (4,8), and (iv) a one-dimensional horizontal edge covering pixels from (7,6) to (8,6).

20 FIG. 6 shows a sample edge map 600 for a simple reference digital image. Edge map 600 shows two edges. The two edges are: (i) a one-dimensional horizontal edge covering pixels from (1,4) to (3,4), and (ii) a one-dimensional horizontal edge covering pixels from (5,3) to (9,3).

Process 400 includes determining descriptive criteria for the edge map for the red digital image (420) and for the edge map for the reference digital image (430). The descriptive criteria may be determined for each edge pixel in the edge or may be determined jointly for multiple edge pixels within an edge (that is, up to and including all 5 edge pixels within the edge). Hereafter the term edge may be used to describe either one edge pixel within an edge or multiple edge pixels within an edge (that is, up to and including all edge pixels within the edge). Descriptive criteria include, for example, (i) whether the edge is a horizontal and/or a vertical edge, (ii) whether the edge transitions, in a given direction, from high to low intensity or low to high intensity—an intensity- 10 change direction, (iii) the location of the edge, (iv) the extent of the edge, (v) the range of intensities that is traversed in the edge, and (vi) various other functions of the pixel intensities and pixel locations of the edges of the red digital image and the reference digital image.

The “edge extent” refers to the set of pixels that define the edge transition. The 15 edge extent can be determined from a given edge using a variety of factors, such as, for example, intensity values. The edge extent of a particular edge also may be influenced by whether there are other edges in the neighborhood of the particular edge. The edge extent may be determined in one or more directions; for example, the edge extent may be determined for either one or two dimensions depending on whether the edge map 20 contains one or two dimensional edges. An edge extent also may be defined for a single edge pixel.

FIG. 7 shows a sample edge extent 700 for the four edges in edge map 500. The edge extent is shown in cross-hatching to differentiate from the actual edge itself,

although the edge is also considered part of the edge extent. As shown, the edge extent for the edge beginning at (1,3) extends upward by 2 pixels and downward by 2 pixels. The edge extent for the edge beginning at (5,1) extends upward by 1 pixel and downward by 3 pixels. The edge extent for the edge beginning at (3,8) extends upward by 1 pixel 5 and downward by 1 pixel. The edge extent for the edge beginning at (7,6) extends upward by 1 pixel and downward by 2 pixels.

FIG. 8 shows a sample edge extent 800 for the two edges in edge map 600. As in FIG. 7, the edge extent is shown in cross-hatching to differentiate from the actual edge itself, although the edge is also considered part of the edge extent. As shown, the edge 10 extents for the two edges extend 1 pixel in both the upward and downward directions.

Process 400 includes determining, for each edge in the edge map for the red digital image, if the edge matches an edge in the edge map for the reference digital image (440). The term “match” is used not only to indicate that the edges correspond to one another with respect to spatial location, but also to indicate that differential resolution 15 exists between the edge of the red digital image and the edge of the reference digital image and that the edge in the red digital image is considered a candidate for modification. Other implementations may determine a “match” using other criteria, such as, for example, by considering only whether edges correspond spatially. The factors used to evaluate edges may provide information relating, for example, to a spatial 20 relationship, to the existence of differential resolution, or to both. The determination of whether edges are matched may be made in a variety of ways, including, for example, using information about each edge individually as well as information comparing edges. In one implementation, descriptive criteria for the edges is compared by considering, for

example, whether the edges have a similar direction, whether the edge extent intensity values transition in the same direction (for example, low to high, or high to low), whether the edge extent intensity values have similar intensity ranges, and whether the edges satisfy a particular distance metric (for example, whether the edges, or some designated part of the edges such as their beginnings, ends, or middles, are within a particular distance of each other).

5        Implementations may use one or more of a variety of other factors to determine whether or not edges match. Such factors include, for example, (i) the distance between the location of the maximum intensity value for the edge under consideration in the red digital image and the location of the maximum or minimum intensity value for the edge under consideration in the reference digital image, (ii) the distance between the location of the minimum intensity value for the edge under consideration in the red digital image and the location of the minimum or maximum intensity value for the edge under consideration in the reference digital image, (iii) when there is another edge (termed an adjacent edge) in the red digital image that is within a particular proximity of the current edge under consideration in the red digital image, the difference between the average intensity of the red image and the reference digital image measured over the edge extent (or a subset of the edge extent) of this adjacent red edge, (iv) the difference between the average intensity of the red edge and the reference edge, as measured over the edge extent (or a subset of the edge extent) of either the red edge or the reference edge, (v) the ratio between the intensity range traversed by the reference edge and by the red edge, and (vi) the difference between the intensity values (or between averages over a range of pixels) of the red and reference edges at various locations. Such factors also may include

frequency information or resolution information in the wavelet domain in addition to spatial domain information. Implementations may determine which factor(s) to use based on the particular characteristics of the red and reference edges. For example, different factors may be used depending upon whether the red and/or reference edges have a 5 maximum intensity value that exceeds a particular threshold value. As another example, different factors may be used depending upon whether one or more edge pixels adjacent to the edge pixel(s) under consideration in the non-reference digital image (red, for example) have matching edge pixel(s) within the reference digital image.

In making the determination of whether two edges match, one or more factors 10 may be analyzed or combined in many ways. For example, the determination may be based on (i) a binary result, such as, for example, whether a particular threshold was met, (ii) a multi-factor analysis, such as, for example, whether a majority of the factors being considered yielded favorable results, where the factors may be weighted in importance, and (iii) a probabilistic determination, such as, for example, a probabilistic assessment of 15 the likelihood that two edges match, where the assessment is based on various factors. In one implementation involving a probabilistic determination, labeling a pixel as a PM pixel may indicate that there is a high probability that the pixel is a matching pixel, but that the probability is lower than that for an M pixel. The preceding examples are not necessarily exclusive of each other.

20 As with determining a match, the determination that a pixel should be labeled as a PM pixel may depend on numerous factors or a single factor. Whether a PM pixel eventually is determined to be an M pixel or an NM pixel may depend upon subsequent

operations. For example, as described below, the continuity, pruning, and relabeling operations may be used to determine how PM pixels are finally to be labeled.

FIGS. 5 and 6 can be used to provide an illustration of operation 440. As discussed above, the twelve edge pixels define four edges in edge map 500. In this example, each edge pixel (and its associated edge extent—in this example, the edge extents run vertically) in edge map 500 is considered in turn to determine if the edge pixel (and its associated extent) matches an edge pixel (and its associated extent) in edge map 600. In other implementations, multiple edge pixels of an edge (up to and including the entire edge) can be jointly considered to determine if the edge matches an edge in edge map 600.

In FIG. 5, edge pixels occur at locations (1,3) to (3,3), locations (5,1) to (9,1), locations (3,8) to (4,8), and locations (7,6) to (8,6). Assuming that a distance of two pixels is acceptable, and that any other matching criteria are satisfied for the edge pixels at locations (1,3), (2,3), (3,3), (5, 1), (6,1), (7,1) and (8,1) when compared to edge pixels (1,4) (2,4), (3,4), (5,3), (6,3), (7,3), and (8,3), respectively, in edge map 600, then these edge pixels are labeled as M pixels. The edge pixel at location (9,1) in edge map 500 is assumed, for purposes of illustration, to be a potential match to edge pixel (9,3) in edge map 600. The edge pixels at locations (3,8) and (4,8) are assumed not to match any edge pixels in edge map 600. The edge pixels at locations (7,6) and (8,6) are assumed, for purposes of illustration, to be a potential match for edge pixels in edge map 600.

Process 400 includes labeling each edge pixel of the edge map for the red digital image (450). The pixels may be labeled, for example, as either an M pixel if there is a

matching reference edge, a PM pixel if there is a potentially matching reference edge, or as an NM pixel otherwise.

FIG. 9 shows an edge map 900 obtained after applying operation 450 to edge map 500. In FIG. 9, the eight pixels at locations (1,3), (2,3), (3,3), (5,1), (6,1), (7,1), and (8,1) 5 are labeled as M pixels because these edge pixels (and their associated edge extents) have matches in edge map 600 for the reference digital image. The pixel at location (9,1) is labeled as a PM pixel. The two pixels at locations (3,8) and (4,8) are labeled as NM pixels, and the two pixels at locations (7,6) and (8,6) are labeled as PM pixels.

Process 400 includes providing continuity among the M and PM pixels in the 10 edge map (460). For example, a neighborhood of a particular extent in the horizontal and/or vertical direction may be defined around each M and PM pixel within the edge map for the lower resolution image. If a pixel within this neighborhood is identified as 15 an M or PM pixel, particular pixels that lie along the direction (for example, horizontal, vertical, or diagonal) between that pixel and the M or PM pixel under consideration may be labeled as M or PM pixels as well in order to ensure continuity among the M or PM pixels. In one such implementation, all PM pixels are treated as M pixels for operation 460 such that all pixels that are relabeled in operation 460 are relabeled as M pixels, and not as PM pixels.

FIG. 10 shows an edge map 1000 obtained after applying operation 460 to edge 20 map 900 using a neighborhood having an extent of five in both the horizontal and vertical directions. With an extent of five in both the horizontal and vertical directions, the neighborhood around pixel (3,3) is shown by a bold outline. Continuity operation 460 changes pixel (4,2) to an M pixel, as shown, because pixel (4,2) lies between M pixel

(5,1) and the M pixel under consideration—pixel (3,3). Edge map 1000 also shows with a bold outline a neighborhood having an extent of five in both the horizontal and vertical directions around pixel (7,6). However, continuity operation 460 does not result in changing any pixels in the neighborhood around pixel (7,6). Pixels (3,8) and (4,8) are not 5 explicitly labeled NM pixels because all unlabeled pixels are understood to be NM pixels.

Continuity between edge pixels may be provided in various other ways. One example includes connecting two edge pixels if the intensity difference between the two edge pixels, as well as the intensity difference between particular pixels that are spatially located between the two edge pixels, both fall below a particular threshold.

10 Process 400 includes removing spurious M or PM pixels in the edge map (470). In one implementation, the set of M or PM pixels are pruned to a smaller set so as to reduce the inclusion of pixels that may have been misclassified as M or PM pixels. Such misclassifications may occur, for example, due to noise. A variety of criteria may be used to perform this pruning operation.

15 FIG. 11 is a flow chart of a process 1100 used by classification unit 330 for pruning M or PM pixels. Process 1100 includes identifying a connected set of M, and optionally PM, pixels (1110). A connected set is defined as a set of all possible M or PM pixels in which each M or PM pixel is adjacent to at least one other M or PM pixel in the set. M or PM pixels are adjacent if they are next to each other in either the horizontal, 20 vertical, or diagonal direction. In edge map 1000, the eight M pixels and one adjacent PM pixel are in a connected set, and the two remaining PM pixels are in a separate connected set.

Process 1100 includes determining, for a given connected set, the number of M pixels (1120) and the number of PM pixels (1125). Process 1100 includes comparing the number of M and/or PM pixels in the connected set to one or more thresholds (1130), and relabeling the M or PM pixels in the connected set to NM pixels if one or more of the 5 thresholds is not satisfied (1140). In one implementation, the sum of the number of M pixels and the number of PM pixels is compared to a size threshold. If the sum is less than the size threshold, then the M and PM pixels are relabeled as NM pixels.

Conversely, if the sum is greater than or equal to the size threshold, then the PM pixels in the connected set are relabeled as M pixels. Operations 1130 and 1140 also, or 10 alternatively, may include comparing the ratio of M to PM pixels in the connected set to a ratio threshold. If the ratio of M to PM pixels is less than the ratio threshold, then the M or PM pixels in the connected set are relabeled as NM pixels. Conversely, if the ratio is greater than or equal to the ratio threshold, then the PM pixels in the connected set are relabeled as M pixels.

15 Continuing with the example above from FIG. 10, and assuming that a size threshold of three is to be applied, then the connected set of M pixels and one PM pixel are not relabeled as NM pixels, and the connected set of PM pixels are relabeled as NM pixels. If a ratio threshold were applied, and assuming that the ratio threshold was less than or equal to 8 M pixels to 1 PM pixel, then the connected set of M pixels would 20 remain as M pixels, the adjacent PM pixel at location (9,1) would be relabeled as an M pixel, and the connected set of PM pixels would be relabeled as NM pixels.

Referring again to FIG. 4, process 400 also includes providing an edge extent for all M pixels in the edge map (480). For M pixels that are edge pixels from the edge map

for the red digital image, the edge extent may be taken to be, for example, the original edge extent from the red digital image. The original edge extent may be determined, for example, in one dimension (for example, horizontal or vertical) or two dimensions. For M pixels that are not edge pixels, such as pixels labeled as M pixels in operation 460, an edge extent may be taken to be, for example, a function of the edge extent of one or more adjacent M pixels.

In one implementation, similarity of edge extents for a connected set of M pixels is provided by designating similar edge extents for all M pixels within the connected set. The edge extents can be determined, for example, based on a function of the original edge extents of the M pixels in the connected set that have an original edge extent. For example, the edge extent for all M pixels in the connected set can be the median, average, minimum, or maximum of the original edge extents. These median, average, or other types of statistical values determined by the collection of M pixels in the connected set also can be used as a starting point in determining the edge extent for each M pixel in the connected set.

FIG. 12 shows an edge extent 1200 obtained by applying operation 480 to the M pixels in edge map 1000. The original edge extent of the M pixels that are also edge pixels is shown in FIG. 7. The edge extent in FIG. 7 has a median upward extent of one pixel and a median downward extent of three pixels. These two medians of the original edge extent are used for the edge extent provided in FIG. 12, and the M pixel that was labeled in operation 460 is also given the same upward and downward extents.

Process 400 includes labeling all edge extent pixels in the edge map as M pixels and labeling all other pixels as NM pixels (490). FIG. 13 shows a classification map

1300 produced by labeling the edge extent pixels of FIG. 12 as M pixels, with the exception that unlabeled pixel locations are understood to be NM pixels.

### Modification Unit

5 Referring again to FIG. 3, system 300 also includes a modification unit 340 that (i) receives classification map  $C_R$  and the three digital images R, G, and B, (ii) increases the resolution of the locations identified by the classification map, and (iii) provides a modified digital color component image for the red digital image (labeled  $M1_R$ ).

Modification unit 340 modifies the resolution content of the red digital image using 10 resolution information from the reference digital image. Because the red digital image and the reference digital image have matching edges, the resolution information from the reference digital image (which is typically a higher resolution digital image compared to the red digital image) can be used to guide the modification of the red digital image.

Modification unit 340 produces resolution information through the application of a 15 wavelet transformation to the reference digital image, and uses the resolution information to modify information produced by the application of a wavelet transformation to the red digital image. A wavelet transformation is useful because such a transformation produces multiple representations (subbands) of the entire image at different resolutions, and because the resolution information at particular spatial locations is correlated with 20 particular wavelet coefficients at corresponding locations within each subband.

FIG. 14 is a flow chart of a process 1400 that modification unit 340 uses to modify the red digital image. Process 1400 includes performing separate wavelet transformations on the red digital image and on the reference digital image (1410). The

transformation may be performed, for example, on either one dimension at a time or two dimensions at a time, and either approach may produce better results in particular implementations. In one implementation that uses two passes, the transformations are one-dimensional in each pass and capture information corresponding to edges having 5 orthogonal directions. The transformation is performed row-wise in the first pass to capture horizontal resolution information. A subsequent pass performs the transformations on columns to capture vertical resolution information.

FIG. 15 is a table 1500 showing the results of a simplified example involving one row. Row “x” refers to the classification map and shows that row x has sixteen pixels, 10 seven of which are M pixels (the intensity values of row x of the digital image are not shown). The seven M pixels are in positions 1-6 and 9, and the remaining pixels are NM pixels, although the NM pixels are not labeled. In general, a one-level, one-dimensional row-wise wavelet transformation of the digital image produces two subbands having a total of sixteen coefficients (not shown) that are correlated or associated with row x of the 15 digital image. One subband represents low resolution information and the other subband represents high resolution information. A four-level, one-dimensional wavelet transformation of the digital image produces five subbands, where there is one coefficient for subband 0 (baseband), one coefficient for subband 1, two coefficients for subband 2, four coefficients for subband 3, and eight coefficients for subband 4. The positions in 20 row x to which each of these sixteen coefficients are correlated are indicated by the numbers shown in table 1500.

In general, the wavelet transformation performed by modification unit 340 on the entire digital image produces, for each row of the digital image, “z” coefficients for the

highest subband, where “z” is equal to half the number of pixels in a row, “z/2” coefficients for the second highest subband, “z/4” coefficients for the third highest subband, and so forth. The wavelet coefficients are said to be correlated or associated with a spatial region (a particular row or portion of a row in this case) of the digital image, although the resolution component identified by the wavelet coefficient will not generally be spatially limited to that region.

5

After computing the wavelet transformations, process 1400 determines which wavelet coefficient locations to modify in the digital image that is produced by the application of the wavelet transformation to the red digital image. This determination is 10 performed for each row separately, and for each subband separately within each row.

Process 1400 includes determining, for a specific row in the red digital image, which wavelet coefficients in the highest subband to modify (1420). In one implementation, the coefficient located at position “j” in the highest subband of the wavelet transformation is modified if position  $2*j$  or  $2*j+1$  in the specific row of the 15 classification map  $C_R$  is an M pixel, where “j” is greater than or equal to zero and is measured from the beginning location of the highest wavelet subband.

Operation 1420 can be illustrated by examining table 1500. The highest subband is subband 4, and position 0 will be modified because position 1 in row x of the classification map is an M pixel. Similarly, positions 1, 2, 3, and 4 will be modified, but 20 positions 5, 6, and 7 will not be modified.

Process 1400 includes determining, for the same specific row in the red digital image, which wavelet coefficients in the remaining subbands to modify (1430). In one implementation, the remaining subbands are processed from highest to lowest. The

coefficient located at position “j” in the next highest subband of the wavelet transformation is modified if position  $2*j$  or  $2*j+1$  in the previous (higher) subband was modified.

Operation 1430 can be illustrated by examining table 1500. Continuing with the 5 previous illustration, the next highest subband is subband 3, and position 0 ( $j=0$ ) will be modified because position 0 ( $2*j$ ) (or position 1 ( $2*j+1$ )) in subband 4 was modified. Similarly, positions 1 and 2 will be modified, but position 3 will not be modified.

Next, subband 2 is processed. Position 0 ( $j=0$ ) will be modified because position 10 0 ( $2*j$ ) (or position 1 ( $2*j+1$ )) in subband 3 was modified. Similarly, position 1 will also be modified.

Next, subband 1 is processed. Position 0 ( $j=0$ ) will be modified because position 0 ( $j=0$ ) (or position 1 ( $2*j+1$ )) in subband 2 was modified.

Finally, subband 0 is processed. Position 0 ( $j=0$ ) will be modified because 15 position 0 ( $j=0$ ) in subband 1 was modified. In one implementation, no coefficients from subband 0 are ever modified.

Process 1400 includes modifying the coefficients in the locations determined above (1440). In one implementation, the coefficients are modified by being replaced. In particular, the coefficients from the determined locations in the result produced by the application of the wavelet transformation to the reference digital image are copied (or 20 scaled values are copied) over the coefficients in the determined locations in the result produced by the application of the wavelet transformation to the red digital image. Thus, the resolution information (or a function of the resolution information) for the reference digital image replaces the resolution information for the red digital image in the

determined coefficient locations. Scaling a coefficient refers to multiplying the coefficient by a particular number.

Note that modification unit 340 only modifies the red (lower resolution) digital image and does not modify the reference (higher resolution) digital image.

5        Process 1400 includes performing an inverse wavelet transformation on the modified result of the application of the wavelet transformation to the red digital image (1450). The inverse transformation produces a modified red digital image that is labeled  $M1_R$  in system 300.

#### Post-Processing Unit

10       Referring again to FIG. 3, system 300 also includes a post-processing unit 350 that receives  $M1_R$ ,  $C_R$ , and the three digital images R, G, and B, and that produces a modified  $M1_R$  referred to as  $M2_R$ . Post-processing unit 350 performs three operations, although other implementations may omit one or more of these three operations and add other operations.

15       First, post-processing unit 350 ensures that pixels in  $M1_R$  that correspond to NM pixels in  $C_R$  were not modified from their intensity values in the red digital image. If any such  $M1_R$  pixels were modified, then the intensity values of those pixels are returned to the original value provided in the red digital image.

NM pixels may have been modified by changing wavelet coefficients in 20 modification unit 340. As described earlier with respect to table 1500, coefficient 0 was modified in subbands 4, 3, 2, 1, and 0. As can be seen in table 1500, these five coefficients all affect NM pixel 0 in row x of the digital image. Thus, NM pixel 0 may be

expected to be modified. Further, although the coefficients are associated or correlated with specific pixels in the digital image, the coefficients typically also produce effects in surrounding pixels (referred to as a spreading effect). Thus, for example, coefficient 3 in subband 4 may be expected to affect not only NM pixel 7 in row x of the digital image, 5 but also NM pixel 8.

Second, post-processing unit 350 computes the resolution of  $M1_R$ , or a portion of  $M1_R$ , to ensure that the resolution did not decrease as a result of the modifications. If the resolution did decrease, a variety of actions may be initiated or performed. For example, if the resolution was decreased at particular pixel locations because of the modifications, 10 then these pixel values may be returned to their original intensity values. Alternatively, one of the parameters for classification unit 330 and/or modification unit 340 may be changed, and the operations of that unit may be performed again. For example, the parameters for determining  $C_R$  may be changed, the reference digital image may be changed, or the way in which wavelet coefficients for modification are selected and/or 15 modified may be changed.

Third, post-processing unit 350 attempts to reduce the perceptual effects of one or more discontinuities that may exist by “feathering” (or, equivalently, smoothing) at the boundaries between M pixels and NM pixels, or between two M pixels that use different reference digital images (Ref1 and Ref2, respectively). Feathering may be applied, for 20 example, and as explained below for several implementations, to locations containing one type of pixel (for example, to pixels on a single side of a transition boundary), or to locations containing multiple types of pixels (for example, to pixels within a neighborhood extending across a transition boundary between different types of pixels).

Feathering is performed, in one implementation, along the vertical boundaries between two horizontally neighboring pixel areas and along the horizontal boundaries between two vertically neighboring pixel areas.

The sizes of the feather extents typically impact the rate at which the intensity values for the pixels at the boundaries blend from one value to another. Various techniques may be used to determine the size of the extents for each row or column corresponding to an NM/M (or Ref1/Ref2) pixel transition. In one implementation, the sizes of the feather extents are determined based on the intensity values of R, G, and B as well as the intensity values of the modified red image,  $M1_R$ . For example, assuming that  $M1_R(i,j)$  is the red intensity value for the NM pixel at the NM/M transition after the red image has been processed by the modification unit, and  $M1_R(i,j+1)$  is the red intensity value for the M pixel at the NM/M transition after the red image has been processed by the modification unit, a value diff may be defined as:

$$\text{diff} = \text{abs}(M1_R(i,j) - M1_R(i,j+1)),$$

which is the absolute value of the red intensity difference between the NM and M red pixel values. In addition, a value diff2 may be defined to be 1 or 0 based on whether a particular metric is satisfied when it is applied to the (R, G, B) pixel intensity values at locations (i,j) and (i,j+1) as well as the  $M1_R$  pixel intensity value at location (i,j+1). For example, the particular metric may be a function of  $\text{abs}(M1_R(i,j+1) - R(i,j+1))$ ,  $\text{abs}(G(i,j) - G(i,j+1))$ , and  $\text{abs}(B(i,j) - B(i,j+1))$ .

An extent can then be defined as:

$$E1 = (\text{constant} * \text{diff}) \text{ if } \text{diff2} = 1, \text{ or } 0 \text{ if } \text{diff2} = 0,$$

where constant  $\geq 0$ , but typically is less than 1.

FIG. 16 provides a simplified example involving a portion of a single column, i.

In FIG. 16, pixel  $(i,j)$  is an NM pixel and  $(i,j+1)$  is an M pixel.  $M1_R(i,j)=130$ ,  $M1_R(i,j+1)=90$ , and  $diff=40$ . Also, assume  $diff2$  is 1. Assuming “constant” has a value of 0.1, then  $E1=4$  which includes  $M1_R(i,j+1)$  through  $M1_R(i,j+4)$ .

5 The extent may be further refined based on the similarity of the adjacent pixel intensities with the same designations. In one implementation, continuing the above example, given an M pixel at location  $(i, j+1)$  and an extent  $E1$ , the following operations are performed for pixels  $(i, k)$ , for  $j+1 < k \leq j+E1$ , where  $k$  is an integer:

10 1. Check whether the pixel  $(i, k)$  is the same designation type as the previous pixel  $(i, k-1)$ , that is, an M pixel or an NM pixel.

2a. If operation 1 is not satisfied, skip to operation 4.

2b. Compare the (R, G, B) intensity values of pixel  $(i,k)$  to the (R, G, B) intensity values of pixel  $(i, k-1)$ .

15 3. If the comparison in operation 2b satisfies a particular metric (which is different from the metric that yields  $diff2$  described above), increment  $k$  and then return to operation 1.

4. If operation 1 is not satisfied or if the comparison in operation 2b does not satisfy a particular metric, the extent is defined as the set of pixels  $(i, k)$  that satisfied operation 3.

20 The above algorithm can be applied to the column shown in FIG. 16. In the first pass through the algorithm,  $k=j+2$ . In operation 1, pixel  $(i,j+2)$  is compared to pixel  $(i,j+1)$ , and both are M pixels. In operation 2b, the (R,G,B) intensity values of pixel  $(i,j+2)$  are compared to the (R,G,B) intensity values of pixel  $(i,j+1)$ . The comparison may

be, for example, to determine the absolute value of a difference for each color component, in which case, the comparison yields a result of (3,3,2) (that is, 128-125, 127-130, 118-116). In operation 3, the result of (3,3,2) is checked to determine if a metric is satisfied. The metric may be, for example, a maximum acceptable difference for each component, and the value of the maximum acceptable difference may be, for example, 6. 5 In such an implementation, the metric is satisfied because the differences for each color component, that is, 3, 3, and 2, are less than the maximum acceptable difference of 6.

In the second pass through the algorithm,  $k=j+3$ . In operation 1, pixel  $(i,j+3)$  is compared to pixel  $(i,j+2)$ , and both are M pixels. In operation 2b, assuming the 10 comparison is the absolute difference in intensity values, the comparison yields a result of 7 (125-118) for the R component and 1 for both the G and B components. In operation 3, assuming the metric is a maximum difference of 5, the result of 7 fails to satisfy the metric. In operation 4, the extent is defined as pixels  $(i,j+1)$  and  $(i,j+2)$ , which is smaller than the earlier determined extent of 4.

15 If the maximum acceptable difference had been 7 in the above example, then pixel  $(i,j+3)$  would have satisfied the metric and a third pass would have been made through the algorithm. In the third pass, however, pixel  $(i,j+4)$  would have failed operation 1 because pixel  $(i,j+4)$  is an NM pixel, so the extent would have ended at pixel  $(i,j+3)$ , including pixels  $(i,j+1)$  through  $(i,j+3)$ .

20 The feather extent can be extended across both M and NM pixel locations, across only M pixel locations, or across only NM pixel locations (or, analogously, across both Ref1 and Ref2 pixel locations, across only Ref1 pixel locations, or across only Ref2 pixel locations in the case where there is an Ref1/Ref2 boundary transition). If the feather

extent is extended across the NM pixel locations as well as the M pixels, an analogous procedure to the operations described above, for example, may be applied to the NM pixels near the NM/M transition boundary.

Once the feather extent is obtained, the intensity values of the pixels within this 5 extent may be modified. In one implementation, within the feather extent, new intensity values are obtained by linearly interpolating between the intensity values associated with the ending pixels of each feather extent. However, many techniques may be used to obtain the new intensity values within the feather extent.

Fig. 17 provides a two-dimensional example 1700 identifying several pixels near 10 the NM/M border transformation that are affected by the feathering scheme described above. For ease of viewing, the NM/M pixels are not labeled. Instead, the transition between the NM/M pixels is indicated by solid bold lines. The pixels affected by the feathering scheme in both the vertical and horizontal directions are labeled as 'S.'

As discussed briefly above, a one-dimensional wavelet transform may have been 15 applied in only one direction by the modification unit 340 using a classification map that described edges in only one direction. In system 300, the unmodified red digital image  $R$  is used to obtain a new classification map in an orthogonal direction. The new classification map and the "old"  $M2_R$  are received by the modification unit 340 which uses the new classification map to apply a one-dimensional wavelet transform in the 20 orthogonal direction to the old  $M2_R$ .  $M2_R$  is then modified to produce a new  $M1_R$ . The new  $M1_R$  is sent to post-processing unit 350 where a final  $M2_R$  is generated. Other implementations may combine the results of multiple passes, orthogonal or otherwise, in various ways.

In another implementation, as with the above, a one-dimensional wavelet transform is applied in only one direction by modification unit 340 using a classification map that describes edges in only one direction. However, in this other implementation,  $M2_R$  is fed back to classification unit 330 where  $M2_R$  is used in place of the red digital image R to obtain a new classification map that indicates M pixels related to edges in an orthogonal direction.

5

### Composite Unit

Referring again to FIG. 3, system 300 also includes a composite unit 360 that receives  $M2_R$  and the two higher resolution digital images, and combines these three digital images to produce a composite color image (frame). An optical printer, as previously described, may be used to combine the three digital images and produce the composite color image. Also, a laser film printer may be used to avoid the loss of resolution that is typically incurred in all three colors with an optical printer.

10

### Additional Implementations

15 Referring again to digitization unit 310, in one implementation,  $I(x,y)$  represents the logarithm of the actual intensity value at pixel location (x,y). Other, implementations may (i) perform one or more of a variety of other transformations, such as, for example, positive-to-negative, in lieu of or in addition to the negative-to-positive transformation, (ii) may perform no transformation at all, (iii) may accept digitized data so as to obviate 20 the need for digitization, and/or (iv) may accept composited data that can be digitized and from which separate digitized component data can then be extracted . For example, (iv)

may include a composite dupe from film separations or a composite color image.

Digitization unit 310 need only digitize two separations if a high resolution separation and a low resolution separation are determined (or designated) beforehand, or determined by digitization unit 310. Further, none of the other blocks in system 300 needs to receive  
5 a digital representation of an unused separation.

Referring again to classification unit 330, implementations may determine the reference digital image by, for example, simply selecting the green or blue digital image, or by selecting the digital image that meets a particular criterion, such as, for example, the digital image that has the highest resolution. Resolution may be determined using a  
10 variety of techniques. For example, frequency information from a transformation may be used to determine the high frequency content of a digital image or separation, or spatial domain techniques may be used to examine edge slopes or other edge features indicative of high frequency or high resolution. Note that resolution determination techniques also may be applied to determine which separation or digital image has the lowest resolution,  
15 although in filming applications it typically may be assumed that the red separation has the lowest resolution.

Additionally, the reference digital image need not be fixed for a given image (frame). For example, the reference digital image may vary depending on the edge or pixel being considered. Further, if it is determined that a particular edge in the lower  
20 resolution image contains a matching edge in more than one of the other digital images, then the reference digital image may be iterated through the two possibilities in multiple passes through classification unit 330 to determine which digital image, or the combination, is a more preferred reference. In one implementation, a set of connected

edge pixels always use the same reference digital image. This reference digital image may be determined based on which reference digital image each of the edge pixels selected, that is, the reference digital image selected by the majority of the edge pixels.

In these examples in which more than one reference digital image can be used

5 within a given image, the classification unit 320 also may specify which one or more of the reference digital images are to provide the information that will be used to modify the selected portions of the lower resolution image. The specification of the reference digital image(s) and the identification of the selected portions to which a given reference digital image applies may be provided in the classification map, as discussed earlier, or

10 elsewhere.

Other implementations determine the low resolution digital image for different portions or features of the frame. In one such implementation, multiple passes are performed through system 300 with each pass processing features in a particular digital image that have the lowest resolution among the three digital images.

15 Various criteria, such as those described above, may be used to determine the resolution of the digital images, or the separations. Such criteria also may include, for example, information obtained from wavelet, or other, transformations. If two digital images have similar resolutions, a reference digital image for a particular edge may be determined between the two by, for example, selecting the digital image with the lower

20 average intensity value at the ending of the edge extent. A digital image may be selected as the reference even if the resolution is not the highest. For example, a digital image may possess some other property that makes the digital image suitable for use as a reference in a particular application.

As stated earlier, classification unit 330 may use another feature in addition to, or in lieu of, an edge to determine which of the pixels to modify. Examples of such features include characteristics of intensity values (for example, intensity values above a particular threshold), object shapes, wavelet coefficient values, areas previously identified as misaligned by some other image processing algorithm, and information from adjacent frames that indicates that a corresponding region in that frame had pixels that were modified (that is, temporally-based features). Temporally-based features may use information from one or more frames that precede, or succeed, or both precede and succeed the frame under consideration in time. Such implementations may take advantage of the fact that much of a frame, including edges and other features, may remain constant from frame to frame.

Alternate implementations of classification unit 330 may perform process 400 using one or more of a variety of edge detection methods in addition to, or in lieu of, the Canny filter mentioned earlier. For example, process 400 may obtain a set of edges by identifying transition pixels at a boundary between higher intensity pixels and lower intensity pixels. Higher intensity pixels may be differentiated from lower intensity pixels by, for example, designating a pixel as higher intensity if the pixel's intensity value is above a threshold, and designating the pixel as lower intensity otherwise. Such a threshold comparison may be used for R, G, and B images.

Various functions performed by classification unit 330, or other units, need not be performed in all implementations. For example, continuity need not be provided among the M or PM pixels in an edge map, and spurious M and PM pixels need not be removed from an edge map. In addition, during the intermediate steps of the various functions

performed by classification unit 330, pixels may be restricted to only the M or NM labels (that is, the PM label need not be assigned to any pixels).

Various functions performed by classification unit 330 may also be performed multiple times in some implementations. For example, continuity and/or pruning steps 5 may be applied multiple times, and at various different stages, during the process performed by classification unit 330. In one alternate implementation, continuity and/or pruning steps may be applied after an initial set of edges (or edge map) is generated but before the various criteria tests are applied to the edges in the non-reference and reference digital images in order to determine whether they match (that is, whether they should be 10 modified).

The slope of intensity-value changes may be used as a descriptive criterion in operations 420 and 430 of process 400. Slope information also may be used in other contexts or for other purposes, such as, for example, to determine edge extent and to determine if edges match (by comparing slope information).

15 Various criteria may also be used to determine whether particular M pixels are to be treated differently than other M pixels during the modification step. When it is determined that particular edges or edge pixels are to be treated differently, this information also may be provided in the classification map or elsewhere. For example, the classification unit may determine different scaling factors to use during the 20 modification operation based on the properties of the edge pixel and its associated extent and the properties of the corresponding edge pixel and its associated extent in the reference digital image.

Referring again to modification unit 340, the resolution content of the images may be modified using time domain analysis, frequency domain analysis, and/or wavelet domain analysis. The implementation described above uses wavelet transformations and may limit the extent of the resolution modifications by ensuring in the time domain that 5 no NM pixels are modified. Other transformations may be used, particularly transformations for which the frequency or resolution information is correlated with spatial or time information, such as, for example, the short-term Fourier transform. Further, different types of transformations may be used within a given implementation. Additionally, temporally-based methods, such as, for example, frame-to-frame analysis, 10 which was mentioned above in the context of classification, may be used to modify an image (frame). Such frame-to-frame analysis may include, for example, many of the techniques already described.

When wavelet (or other) transformations are used, the coefficients may be combined in myriad ways, such as, for example, by copying as described above, or by 15 performing scaling or other functions. Further, the coefficients may be modified based on other factors of the subbands. Additionally, given the spreading effect of a coefficient, one implementation purposefully changes coefficients that are not associated or correlated with a particular M pixel but that are expected to impact the particular M pixel through the spreading effect. Implementations may perform transformations on 20 subsets of the digital images, as opposed to performing transformations on the entire digital image.

The wavelet transformation used in one implementation of modification unit 340 is a digital wavelet transformation that uses subsampling. Other implementations of wavelet transformations, or other transformations, may be used.

Referring again to post-processing unit 350, various other operations may be 5 performed to verify or improve the resolution obtained. For example, post-processing unit 350 may allow NM pixels to be modified and/or may ensure that M pixels were modified in a beneficial manner. In contrast, it also should be clear that implementations need not perform any post-processing.

Implementations need not process every edge in the red digital image. For 10 example, one implementation only attempts to modify edges in the red digital image that have a resolution below a specified threshold or that have a resolution that is lower than the corresponding edge in the reference digital image by a specified threshold.

The implementations and techniques described herein can be applied to a variety of applications in which distortion that results from multiple separations of differential 15 resolution needs to be reduced. Examples include spectral and non-spectral separations. Spectral separations are used, for example, in: (1) color film applications capturing, for example, different color frequencies, (2) astronomical applications capturing, for example, radio frequencies and/or optical frequencies, and (3) medical applications capturing, for example, different magnetic (MRI), X-ray, and sound (ultrasound) 20 frequencies. As these examples illustrate, spectral separations may be captured from various frequency sources, including, for example, electromagnetic and sound waves. Non-spectral separations may be obtained from, for example, variations in pressure, temperature, energy, or power.

The implementations and techniques described herein also may be applied to composite color images, that is, to images that have more than one color component. For example, a video image may have red, green, and blue components combined into one “composite” image. These components may be separated to form separations, and one or 5 more of the implementations and techniques described herein may be applied to the separations.

Implementations and features may be implemented in a process, a device, or a combination of devices. Such a device may include, for example, a computer or other processing device capable of processing instructions using, for example, a processor, a 10 programmable logic device, an application specific integrated circuit, or a controller chip. Instructions may be in the format of, for example, software or firmware. Instructions may be stored in a computer readable medium, such as, for example, a disk, a random-access memory, or a read-only memory.

Referring to FIG. 18, a system 1800 for implementing various disclosed features 15 includes a processing device 1810 coupled to a computer readable medium 1820. Computer readable medium 1820 stores instructions 1830 to be processed by processing device 1810, wherein such processing implements the various disclosed features.

A separation or digital image may be selected, for example, by being accessed. Selecting a separation or digital image may be done, for example, by selecting a file or a 20 representation, such as, for example, a display, of the separation or digital image. Other representations may be provided, for example, by various user interfaces.

The digital images described above include information that generally spans the same object. For example, the red, blue, and green digital images each contain

information (red information, blue information, or green information) that spans the entire frame. Such information (red, blue, and green) can be termed “complementary” information because the information relates to the same object. Other implementations may use digital images that span a different object, such as, for example, an area in a 5 scene being filmed, a portion of a body, or a portion of the sky.

Other implementations may use digital images in which only a portion of each digital image spans the desired object. For example, a red separation might be used that only captures foreground information, and a reference separation might be used that captures both foreground and background information.

10        Implementations also may use digital images in which the desired object is captured from a different angle or distance. Such a situation may occur when cameras at different distances or angles are used to film the different color components, or different telescopes are used to photograph the same general portion of the sky.

15        Various implementations perform, for example, one or more operations, functions, or features automatically. Automatic refers to being performed substantially without human intervention, that is, in a substantially non-interactive manner. Examples of automatic processes include a process that is started by a human operator and then runs by itself. Automatic implementations may use, for example, electronic, optical, mechanical, or other technologies.

20        The functional blocks, operations, and other disclosed features may be combined and performed in different orders and combinations, and may be augmented with other features not explicitly disclosed. Reference to a portion of an image or other object may include the entire image or other object.

A number of implementations have been described. Nevertheless, it will be understood that various modifications may be made without departing from the spirit and scope of the claims. Accordingly, other implementations are within the scope of the following claims.